S-Link to Gigabit Ethernet Adapter
New Frame Segmentation for LHCb Data Acquisition System

Joaquim E. Neves
EP-Division, CERN, CH-1211 Geneva 23, Switzerland.
Escola de Engenharia da Universidade do Minho, P-4800-058 Guimares, Portugal.
Richard Jacobsson, Niko Neufeld, and Beat Jost
EP-Division, CERN, CH-1211 Geneva 23, Switzerland

arXiv:physics/0304014v1 [physics.ins-det] 4 Apr 2003

Data Acquisition and Control Systems used in high energy physics experiments, such as those which will take place in the Large Hadron Collider (LHC) at CERN, require the specification of data formats and transmission protocols as well as the use of high speed links and interfaces.
In this context, a new Frame Segmentation process will be presented and discussed, based on data formats adopted by the LHCb experience for the interconnection of two standardized systems: S-link and Gigabit Ethernet. Simulation results of the transfer capacity of the proposed mechanism will be also reported, together with guidelines for its physical implementation.

1. INTRODUCTION

The Large Hadron Collider beauty experiment for precise measurements of CP violation [1] and rare decays (LHCb) [2] requires High Speed Interconnect (HSI) Systems in order to transport the large amount of data generated by the detectors connected to the Front-end Motherboards (FEMB) to the storage devices connected to the Read-out Motherboards (ROMB).

In the opposite direction, only a small volume of data generally has to be transferred between those parts for the control and management of the whole system.

The LHCb experiment adopted Gigabit Ethernet (GbE) [3] as the link technology from the output of the FEMB electronics boards to the input of the ROMB Sub-Farm Controllers (SFCs), and S-Link [4], as a standard interface between the FEMB and the Data Acquisition (DAQ) system [2].

The Architecture of the LHCb Data Acquisition, based both on full S-link transmission and on the S-Link to Gigabit Ethernet Adapter (SGbEA) System, is depicted in Figure 1. Since the Readout Network already supports the GbE, one S-Link card for the FEMB, based on GbE, was designed and built within the scope of the collaboration with the Atlas LHC experiment [5].

Next, in section 1, the technical specification of the LHCb Data Acquisition System is discussed, while section 3 describes the new frame segmentation mechanism of the SGbEA. Simulation results for SGbEA throughput, using two different operation modes and data packet lengths, are presented in section 4. The Section 5 deals with the physical implementation of the New Frame Segmentation and Section 6 summarizes the main conclusions.

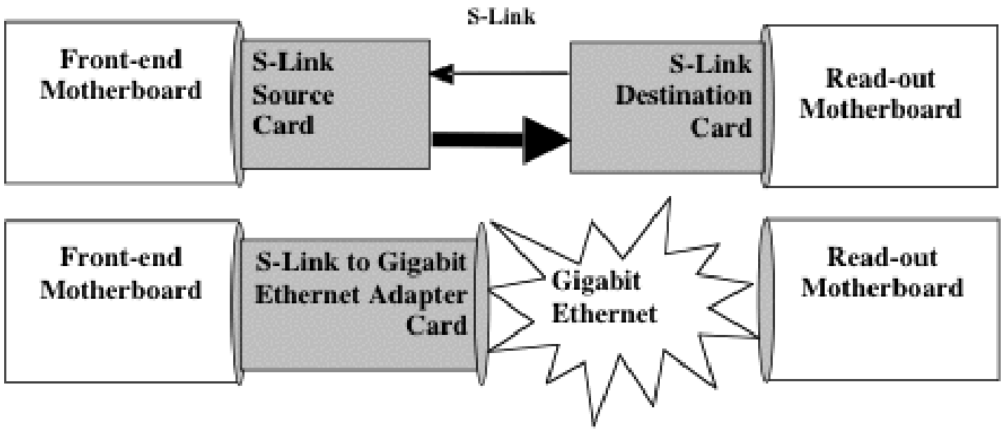

Figure 1: - Architecture of the LHCb Data Acquisition System, based on S-Link to Gigabit Ethernet Adapter.

2. SYSTEM ARCHITECTURE

The S-LINK specification defines a simple FIFO-like user interface at both ends of the transmission link which remains independent of the technology used to implement the physical link and which provides the transfer of event data and control words, error detection, optional flow control and test facilities [6]. The S-link specification also describes the interfaces between the FEMB and the Link Source Card (LSC) and between the Link Destination Card (LDC) and the ROMB, in either simplex or duplex version.

In the simplex version, since there is no communication path from the LDC to the LSC, the transmission is unidirectional, while, in the duplex version, the return channel between the LDC and the LSC allows the transmission of flow control commands from the ROMB to the FEMB. In both versions a single, high-density, 64-pin connector is used to connect the FEMB to the LSC and the LDC to the ROMB, allowing, per clock cycle, the transport of a 32 bit word at the frequency of 40 MHz, which corresponds to a throughput of 1.28 Gbit/s.

There is more than one solution for the implementation of the Front-End Multiplexers (FEMs) and the Readout Units (RUs). This is because, in the LHCb experiment, the main difference between these modules results from the fact that the RUs have to interface to the Readout Network, and hence must respect the GbE flow control protocol, which is not true for the links between FEMs and RUs [2].

The SGbEA is a possible solution to interconnect FEMs and RUs, as is shown in Figure 1. This module has to implement the Ethernet functionalities of the Physical and Media Access Control (MAC) Layers, and the S-link specifications. The protocol conversion between S-link and GbE can be implemented within a FPGA, together with memories (FIFO, RAM and Registers), as can be seen in Figure 2.

The SGbEA has to generate the Start Of Packet (SOP) and the End Of Packet (EOP) signals for the MAC device and to stop and restart the transmission when receiving watermark flags from the MAC or S-link FIFOs. Optionally, it can process the fragmentation of the S-link Packets on Ethernet frames, as is explained in the next section.

3. NEW FRAME SEGMENTATION

Since within the communication links, between LHCb FEMs and RUs, the data can be transported with two different formats, the SGbEA has to support two operational modes: Short Packet Mode (SPM) and Long Packet Mode (LPM).

In SPM, the length of the packets generated by the FEMB is variable, but always smaller than the maximum length of the Ethernet MAC frames. On the other hand, since the header of the LHCb data format already includes the header of the MAC Frames, no protocol conversion is necessary and the S-link packet is forwarded directly from the S-link FIFO to the MAC FIFO.

In LPM, the FEMB can generate data packets ranging in length from 52 Bytes to 32 Kbytes.

In this operational mode the length of the S-link packet can be higher than the maximum length of the Ethernet MAC frames. For this reason, the header of each S-link packet, which contains a field with its length, has to be memorized in the SGbEA RAM, before calculating the number of fragments into which it must be split in order to be transmitted.

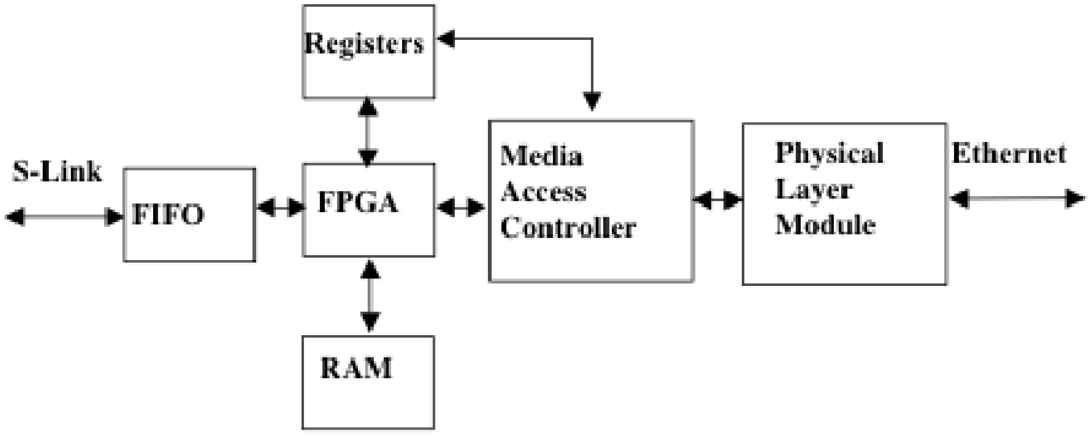

Figure 2: - S-Link to Gigabit Ethernet Adapter Architecture.

In this mode, the SGbEA inserts, within each fragment header, the Type/length of the frame on the Ethernet header field, together with the number of fragments and the current fragment number of the S-link packet.

Optionally, in both operational modes, the Ethernet Sources and Destination Address can be generated by the FEMB (on a packet-by-packet basis), or can be previously inserted in the SGbEA Registers, by the control system and then transferred to each packet fragment header.

4. SIMULATION RESULTS

A VHSIC Hardware Description Language (VHDL) model of the SGbEA module was developed for simulation purposes on the VisualHDL platform on a Computer-aided Engineering (CAE) system, together with an S-link packet generator.

Packets with different lengths and data formats were generated by this model at the S-link interface in either Long or Short Packet Mode, in order to be processed by the SGbEA and transmitted over the Gigabit Ethernet.

In both cases the S-link overhead, introduced by the control words used for signaling the beginning and the end of the LHCb data packets, was reduced to a minimum: one word for the start and another one for the end of packet. In another interface, at the Ethernet side, in addition to the MAC overhead, a minimum inter frame gap of 12 bytes was also guaranteed.

4.1. Short Packet Mode

Since in SPM there is no frame segmentation, the inclusion of a MAC frame header at the beginning of the packet is the only constraint imposed on the data format.

Figures 3 and 4 report some results for the SGbEA simulation in SPM. The maximum throughput for different packet lengths is present in Figure 3, while Figure 4 shows the relationship between the effective throughput reached with those packets, and Gigabit Ethernet capacity.

As the figures show, due to the overhead of the Ethernet frames, when short packets with a size to the order of tens of bytes are transferred, the throughput is clearly lower than the transmission capacity, both in the S-link and in the GbE interfaces. On the other hand, for longer packets with over a hundred bytes, the transfer rate begins to be limited by the effective GbE capacity.

It is interesting to note that, for 64 byte packets, the reference value for packet length on several interfaces of the LHCb data acquisition system, the throughput is near 1.4 M Packet/s, allowing the use of SGBE on these interfaces with a trigger rate of about 1.1 MHz.

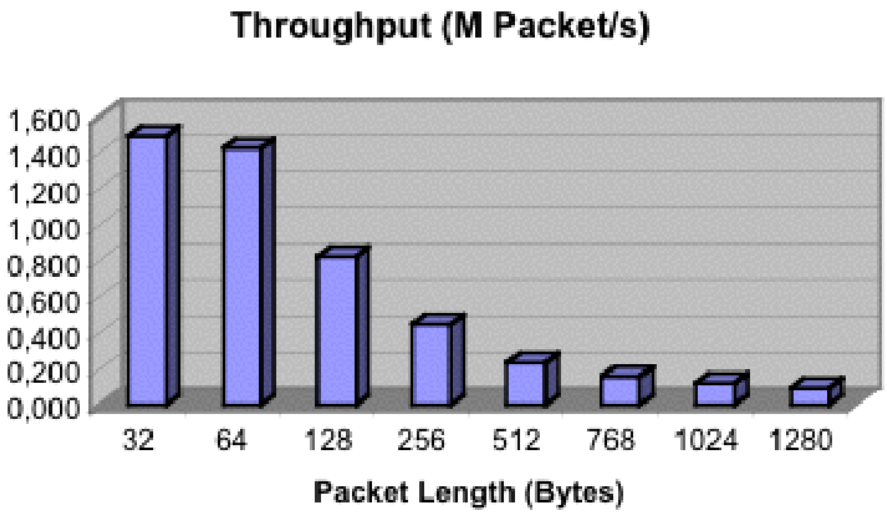

Figure 3: - Maximum throughput achieved by the SGbEA with Short Packet Mode, for different packet length.

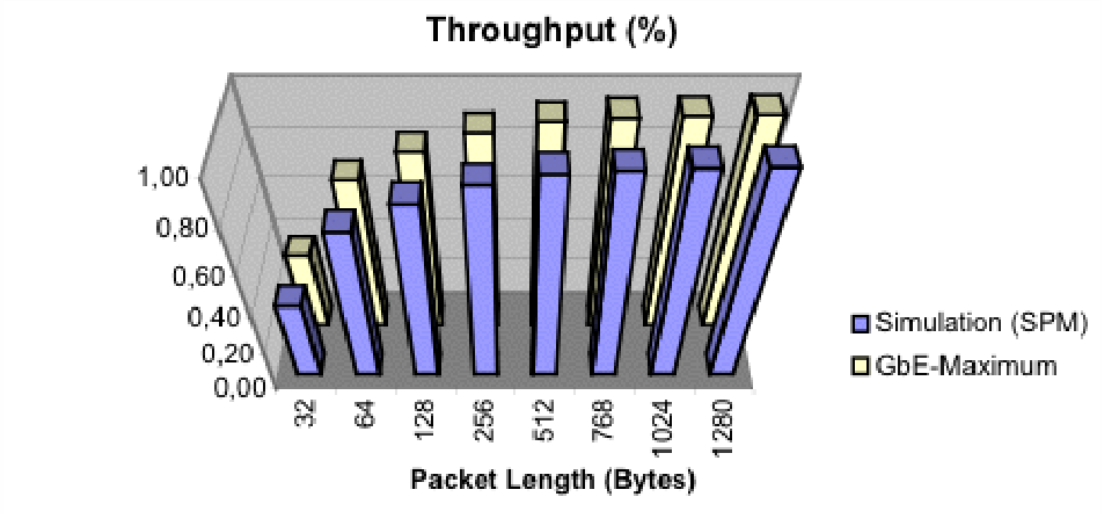

Figure 4: - Maximum throughput achieved by Short Packet Mode and Gigabit Ethernet capacity for different packet length.

4.2. Long Packet Mode

Figures 5 and 6 present simulation results for Frame Segmentation in LPM. The maximum throughput of the system is shown in Figure 5, while Figure 6 shows the relationship between the effective throughput achieved on the transmission of those packets and Gigabit Ethernet capacity.

In this mode, the LHCb data packets presented on the S-link interface are of variable length, as in SPM, but now the packet header has 52 bytes, including the header of the MAC frame.

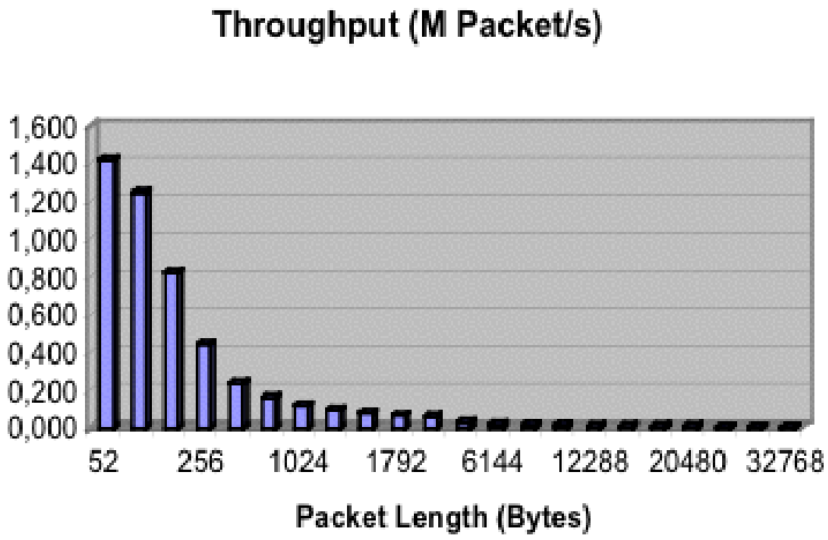

Figure 5: - Maximum throughput achieved by SGbEA in Long Packet Mode, for different packet lengths.

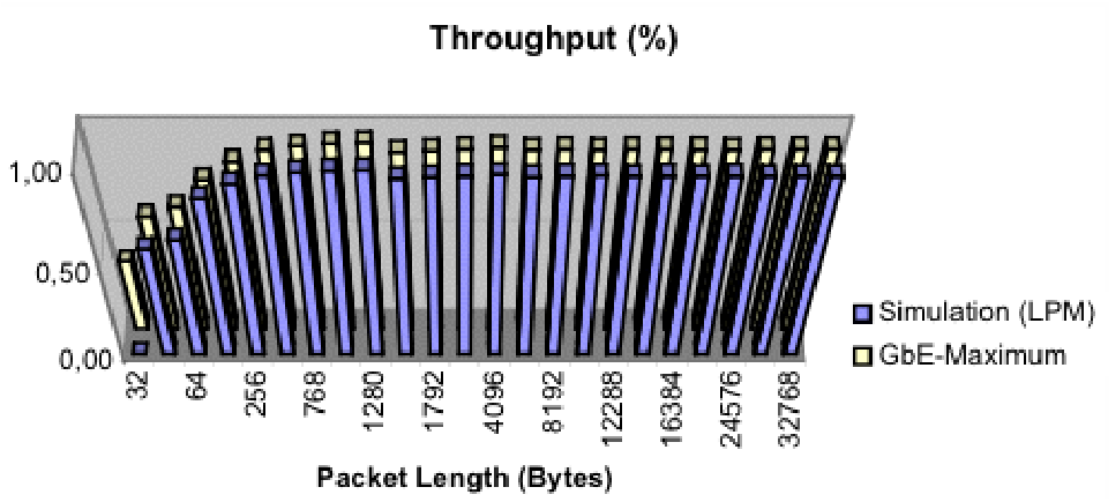

Figure 6: - Maximum throughput achieved in Long Packet Mode and Gigabit Ethernet capacity for different packet lengths.

As was expected, Figure 5 shows that the throughput achieved for packets of a length of less than 64 bytes is lower than that achieved in SPM, but is still higher than 1,2 M Packet/s, which means that the LPM can also be used with these packet lengths on the interfaces with a trigger rate of about 1,1 MHz.

For longer packets, the decrease of the throughput is proportional to the increase of packet length, as the overhead introduced by the fragmentation process becomes negligible. As is depicted in Figure 6, the effective occupancy of the Ethernet payload increases with the length of the packet for short length packets - as was observed in the SPM simulations - and remains constant with longer packets.

5. PHYSICAL IMPLEMENTATION

The Frame Segmentation mechanism for LHCb Data Acquisition System described previously has been implemented within a FPGA on the S-Link to Gigabit Ethernet Adapter board itself, developed at the Argonne National Laboratory for Atlas LHC Experience. This prototyping board was implemented on a small daughter board over a PCI [1] Mezzanine Card (PMC), according to the IEEE Common Mezzanine Card standard [7].

A new version of the SGbEA has been specified for implementation in the near future. Since this new board will already be connected to the FEMB, the new Frame Segmentation can be fully tested, at the maximum S-link throughput, without the present constrains of the PCI Bus.

The single port MAC Controller [8] could be replaced by a dual port [9] that supports the Gigabit Media Independent Interface (GMII) and the Ten-Bit Interfaces (TBI), allowing the Ethernet transmission over copper and/or fiber media. The Physical Layer module could also be implemented with a single dual port transceiver, either for copper or fiber media [10]. While one serial port is used in the copper interface, by connecting this port to a magnetic device [11], the other is used in the fiber interface, by connecting this port to the optical module.

6. CONCLUSIONS

A new Frame Segmentation mechanism, with flow control capabilities, has been implemented over the S-Link to Gigabit Ethernet Adapter for the LHCb Data Acquisition System.

This mechanism is not only able to halt the transmission, in case of overflow, but can also optionally insert the Ethernet Source and Destination Addresses in the MAC frames.

As the simulation results show, this technique permits the transmission of the Short and/or Long Packets onto High Speed Interconnect Systems, which will be used in LHC experiments at CERN.

[1] PCI, Peripheral Component Interface.

References

[1] T. Nakada, CP Violation, CERN Academic Training lecture series, Geneva, Feb 97. http://lhcb-doc.web.cern.ch/lhcb-doc/presentations/Bphysics%20and%20CP%20violation/lecture/cpviol.htm

[2] LHCb Online System - Data acquisition and experiment control, Technical Design Report, CERN/LHCC 2001-40, December 2001. http://lhcb-comp.web.cern.ch/lhcb-comp/General/Publications/odin_tdr.pdf

[3] IEEE Std Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications, IEEE - Std Association, Piscataway, NJ 08855-1331, USA, 2002. http://shop.ieee.org/store/product.asp?prodno=SS94973

[4] O. Boyle et al., The S-LINK Interface Specification, ECP-Division, Geneva, March 97, http://hsi.web.cern.ch/HSI/s-link/spec/spec

[5] Nick Ellis and John Dawson, private communications.

[6] H. C. van der Bij et al., S-LINK, a Data Link Interface Specification for the LHC Era, IEEE Transactions on Nuclear Science, Volume: 44 Issue: 3 Part: 1, Anaheim, CA, USA, June 97.

[7] IEEE P1386 Standard Mechanics for a Common Mezzanine Card Family: CMC, Institute of Electrical and Electronics Engineers, Piscataway, NJ 08855-1331, USA

[8] http://206.204.107.130/products/networking/controllers.html

[9] http://www.intel.com/design/network/products/ethernet/linecard_ec.htm

[10] http://www.marvell.com/Internet/Products/products/1,2414,3-15-151-22,00.html

[11] http://www.trans-power.com/gigganged.html